Compositional Reasoning for Interval Markov Decision Processes

Vahid Hashemi[a,b], Holger Hermanns[b], Andrea Turrini[c]

[a]*Max Planck Institute for Informatics, Saarbrücken, Germany*
[b]*Department of Computer Science, Saarland University, Saarbrücken, Germany*
[c]*State Key Laboratory of Computer Science, ISCAS, Beijing, China*

arXiv:1607.08484v1 [cs.LO] 28 Jul 2016

Abstract

Model checking probabilistic CTL properties of Markov decision processes with convex uncertainties has been recently investigated by Puggelli et al. Such model checking algorithms typically suffer from the state space explosion. In this paper, we address probabilistic bisimulation to reduce the size of such an MDP while preserving the probabilistic CTL properties it satisfies. In particular, we discuss the key ingredients to build up the operations of parallel composition for composing interval MDP components at run-time. More precisely, we investigate how the parallel composition operator for interval MDPs can be defined so as to arrive at a congruence closure. As a result, we show that probabilistic bisimulation for interval MDPs is congruence with respect to two facets of parallelism, namely synchronous product and interleaving.

Keywords: Markov Decision Process, Interval MDP, Compositionality, Probabilistic Bisimulation

1. Introduction

Probability, nondeterminism, and uncertainty are three core aspects of real systems. *Probability* arises when a system, performing an action, is able to reach more than one state and we can estimate the proportion between reaching each of such states: probability can model both specific system choices (such as flipping a coin, commonly used in randomized distributed algorithms) and general system properties (such as message loss probabilities when sending a message over a wireless medium). *Nondeterminism* represents behaviors that we can not or we do not want to attach a precise (possibly probabilistic) outcome to. This might reflect the concurrent execution of several components at unknown (relative) speeds or behaviors we keep undetermined for simplifying the system or allowing for different implementations. *Uncertainty* relates to the fact that not all system parameters may be known exactly, including exact probability values.

Probabilistic automata (*PA*s) [20] extend classical concurrency models in a simple yet conservative fashion. In probabilistic automata, concurrent processes may perform probabilistic experiments inside a transition. *PA*s are akin to *Markov decision processes* (*MDP*s), their fundamental beauty can be paired with powerful model checking techniques, as implemented for instance in the PRISM tool [18].

In *PA*s and *MDP*s, probability values need to be specified precisely. This is often an impediment to their applicability to real systems. Instead it appears more viable to specify ranges of probabilities, so as to reflect the uncertainty in these values. This leads to a model where intervals of probability values replace probabilities. This is the model studied in this paper, we call it *interval Markov decision processes*, *IMDP*s.

In standard concurrency theory, *bisimulation* plays a central role as the undisputed reference for distinguishing the behaviour of systems. Besides for distinguishing systems, bisimulation relations conceptually allow us to reduce the size of a behaviour representation without changing its properties (i.e., with respect to logic formulae the representation satisfies). This is particularly useful to alleviate the state explosion problem notoriously encountered in model checking. If the bisimulation is a congruence with respect to a parallel composition operator used to build up the model out of smaller ones, this can give rise to a compositional strategy to associate a small model to a large system without intermediate state space explosion. In several related settings, this strategy has been proven very effective [5, 11].

Markov chains are known to be closed under interleaving parallelism (if considering the continuous-time setting) and under synchronous (also called synchronous product) parallelism (if considering the discrete-time setting). The more general concept of asynchronous parallelism with synchronisation (as in CCS or CSP) is known to require nondeterminism so as to arrive at closure properties (yielding *PA* for discrete time and interactive *MC* [10] for continuous time).

These observations are conceptually echoed in the setting considered in the present paper, albeit for very different reasons. While nondeterminism is a genuine asset of *IMDP*s, a closure property can not be established for asynchronous parallelism with synchronisation. It has been recently investigated in [9] the possibility of establishing a asynchronous parallelism with synchronisation for *IMDP* models. However, the underlying construction is problematic since it does not manage correctly the spurious distributions. More precisely, for a pair of *IMDP* components the equality of the emerged sets of spurious distributions as a parallelism result should be guaranteed in order to establish the congruence result. This fact is not treated precisely in the setting of [9] for the defined asynchronous parallelism

with synchronisation. In this work instead *IMDPs* are shown to be closed under interleaving parallelism, as well as under synchronous parallelism. This enables us to develop compositionality results with respect to bisimulation for these two facets of parallelism.

Related work. Compositional specification of uncertain stochastic systems has been explored in various works before. Interval *MCs* [13, 17] and Abstract *PAs* [6] serve as specification theories for *MCs* and *PAs* featuring satisfaction relation, and various refinement relations. In order to be closed under parallel composition, Abstract *PAs* allow general polynomial constraints on probabilities instead of interval bounds. Since for Interval *MCs* it is not possible to explicitly construct parallel composition, the problem of whether there is a common implementation of a set of Interval *MCs* is addressed instead [7]. To the contrary, interval bounds on *rates* of outgoing transitions work well with parallel composition in the continuous-time setting of Abstract Interactive *MCs* [16]. The reason is that unlike probabilities, rates do not need to sum up to 1. Authors of [24] successfully define parallel composition for interval models by separating synchronizing transitions from the transitions with uncertain probabilities.

Organization of the paper. We start with necessary preliminaries in Section 2. In Section 3, we give the definition of probabilistic bisimulation for *IMDPs* and discuss the main results of [8]. Furthermore, we show that the probabilistic bisimulation over *IMDPs* is compositional and transitive. Finally, in Section 6 we conclude the paper.

2. Preliminaries

Given $n \in \mathbb{N}$, we denote by $\vec{1} \in \mathbb{R}^n$ the unit vector and by $\vec{1}^T$ its transpose. In the sequel, the comparison between vectors is element-wise and all vectors are column ones unless otherwise stated. For a given set $P \subseteq \mathbb{R}^n$, we denote by $\mathrm{CH}(P)$ the convex hull of P and by $\mathtt{Ext}(P)$ the set of extreme points of P. If P is a polytope in $\mathbb{R}^n$ then for each $i \in \{1, \ldots, n\}$, the projection $\mathrm{proj}_{\mathbf{e_i}} P$ of P is defined as the interval $[\min_i P, \max_i P]$ where $\min_i P = \min\{ x_i \mid (x_1, \ldots, x_i, \ldots, x_n) \in P \}$ and $\max_i P = \max\{ x_i \mid (x_1, \ldots, x_i, \ldots, x_n) \in P \}$.

We denote by $\mathbb{I}$ is a set of closed subintervals of $[0, 1]$ and, for a given $[a, b] \in \mathbb{I}$, we let $\inf[a, b] = a$ and $\sup[a, b] = b$.

Given a set X, we denote by $\mathcal{I}_X$ the identity equivalence relation $\mathcal{I}_X = \{ (x, x) \mid x \in X \}$. We may drop the subscript X from $\mathcal{I}_X$ when the set X is clear from the context.

Given two relations $\mathcal{R} \subseteq X \times Y$ and $\mathcal{S} \subseteq U \times V$, we denote by $\mathcal{R} \times \mathcal{S}$ the relation $\mathcal{R} \times \mathcal{S} = \{ ((x, u), (y, v)) \in (X \times Y) \times (U \times V) \mid (x, y) \in \mathcal{R}, (u, v) \in \mathcal{S} \}$. If $\mathcal{X}$ is an equivalence relation on X and $\mathcal{Y}$ an equivalence relation on Y, then $\mathcal{X} \times \mathcal{Y}$ is an equivalence relation on $X \times Y$.

For a given set X, we denote by $\Delta(X)$ the set of discrete probability distributions over X and by $\delta_x \in \Delta(X)$ the *Dirac* distribution on x, that is, the distribution such that for each $y \in X$, $\delta_x(y) = 1$ if $y = x$, 0 otherwise. Given two sets X and Y and two distributions $\rho_X \in \Delta(X)$ and $\rho_Y \in \Delta(Y)$, we denote by $\rho_X \times \rho_Y$ the distribution $\rho_X \times \rho_Y \in \Delta(X \times Y)$ such that for each $(x, y) \in X \times Y$, $(\rho_X \times \rho_Y)(x, y) = \rho_X(x) \cdot \rho_Y(y)$. Given a finite set of indexes I, a multiset of distributions $\{ \rho_i \in \Delta(X) \mid i \in I \}$, and a multiset of real values $\{ p_i \in \mathbb{R}_{\geq 0} \mid i \in I \}$, we say that ρ is the convex combination of $\{ \rho_i \in \Delta(X) \mid i \in I \}$ according to $\{ p_i \in \mathbb{R}_{\geq 0} \mid i \in I \}$, denoted by $\rho = \sum_{i \in I} p_i \cdot \rho_i$, if $\sum_{i \in I} p_i = 1$ and for each $x \in X$, $\rho(x) = \sum_{i \in I} p_i \cdot \rho_i(x)$. For an equivalence relation $\mathcal{R}$ on X and $\rho_1, \rho_2 \in \Delta(X)$, we write $\rho_1 \; \mathcal{L}(\mathcal{R}) \; \rho_2$ if for each $C \in X/\mathcal{R}$, it holds that $\rho_1(C) = \rho_2(C)$. By abuse of notation, we extend $\mathcal{L}(\mathcal{R})$ to distributions over $X/\mathcal{R}$, i.e., for $\rho_1, \rho_2 \in \Delta(X/\mathcal{R})$, we write $\rho_1 \; \mathcal{L}(\mathcal{R}) \; \rho_2$ if for each $C \in X/\mathcal{R}$, it holds that $\rho_1(C) = \rho_2(C)$.

2.1. Interval Markov Decision Processes

Let us formally define Interval Markov Decision Processes.

Definition 1. An *Interval Markov Decision Process* (*IMDP*) $\mathfrak{I}$ is a tuple $\mathfrak{I} = (S, \bar{s}, \mathcal{A}, \mathtt{AP}, L, I)$, where S is a finite set of *states*, $\bar{s} \in S$ is the *initial state*, $\mathcal{A}$ is a finite set of *actions*, $\mathtt{AP}$ is a finite set of *atomic propositions*, $L \colon S \to 2^{\mathtt{AP}}$ is a *labelling function*, and $I \colon S \times \mathcal{A} \times S \to \mathbb{I}$ is an *interval transition probability function* such that for each s, there exist a and s' such that $I(s, a, s') \neq [0, 0]$. We denote by $[\mathfrak{I}]$, the class of all finite-state finite-transition *IMDPs*.

We denote by $\mathcal{A}(s)$ the set of actions that are enabled from state s, i.e., $\mathcal{A}(s) = \{ a \in \mathcal{A} \mid \exists s' \in S : I(s, a, s') \neq [0, 0] \}$. Furthermore, for each state s and action $a \in \mathcal{A}(s)$, we let $s \xrightarrow{a} \mu_s$ mean that $\mu_s \in \Delta(S)$ is a *feasible distribution*, i.e., for each state s' we have $\mu_s(s') \in I(s, a, s')$. We require that the set $\mathcal{P}^{s,a} = \{ \mu_s \mid s \xrightarrow{a} \mu_s \}$ is non-empty for each state s and action $a \in \mathcal{A}(s)$.

An *IMDP* is initiated in some state s_1 and then moves in discrete steps from state to state forming an infinite path $s_1 \, s_2 \, s_3 \ldots$. One step, say from state s_i, is performed as follows. First, an action $a \in \mathcal{A}(s)$ is chosen nondeterministically by *scheduler*. Then, *nature* resolves the uncertainty and chooses nondeterministically one corresponding feasible distribution $\mu_{s_i} \in \mathcal{P}^{s_i,a}$. Finally, the next state s_{i+1} is chosen randomly according to the distribution μ_{s_i}. For a more formal treatment of the *IMDP* semantics, we refer the reader to [8, 9].

Observe that the scheduler does not choose an action but a *distribution* over actions. It is well-known [20] that such randomization brings more power in the context of bisimulations. Note that for nature this is not the case, since $\mathcal{P}^{s,a}$ is closed under convex combinations, thus nature can choose all distributions.

2.2. Action Agnostic Probabilistic Automata

We now introduce the action agnostic probabilistic automata we use in this paper, based on the probabilistic automata framework [20], following the notation of [21]. Note that the probabilistic automata we consider here correspond to the *simple* probabilistic automata of [20]. In practice, we consider the subclass of (simple) probabilistic automata of [20] having as set of actions the same singleton $\{f\}$, that is, all transitions are labelled by the same external action f. Since this action is unique, we just drop it from the definitions.

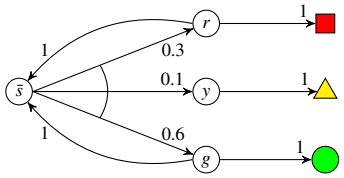

Figure 1: An example of *PAs*: the *PA* $\mathfrak{E}$

Definition 2. An *(action agnostic) probabilistic automaton* (*PA*) is a tuple $\mathfrak{P} = (S, \bar{s}, \mathtt{AP}, L, T)$, where S is a set of *states*, $\bar{s} \in S$ is the *start* state, $\mathtt{AP}$ is a finite set of *atomic propositions*, $L\colon S \to 2^{\mathtt{AP}}$ is a *labelling function*, and $T \subseteq S \times \Delta(S)$ is a *probabilistic transition relation*.

We denote by $[\mathfrak{P}]$ the class of all finite-state finite-transition probabilistic automata and we assume that each state in S is reachable from $\bar{s}$. We may drop *action agnostic* since this is the only type of probabilistic automata we consider. The start state is also called the *initial* state; we let s, t, u, v, and their variants with indices range over S.

We denote the generic elements of a probabilistic automaton $\mathfrak{P}$ by S, $\bar{s}$, $\mathtt{AP}$, L, T, and we propagate primes and indices when necessary. Thus, for example, the probabilistic automaton $\mathfrak{P}'_i$ has states S'_i, start state $\bar{s}'_i$, and transition relation T'_i.

A transition $tr = (s, \mu) \in T$, also written $s \longrightarrow \mu$, is said to *leave* from state s and to *lead* to the measure μ. We denote by $src(tr)$ the *source* state s and by $trg(tr)$ the *target* measure μ, also denoted by μ_{tr}. We also say that s enables the transition (s, μ) and that (s, μ) is enabled from s.

Example 1. An example of *PA* is the one shown in Figure 1: the set of states is $S = \{\bar{s}, r, y, g, \blacksquare, \blacktriangle, \bullet\}$, the start state is $\bar{s}$, the set of atomic propositions is $\mathtt{AP} = S$, the labelling function L is such that for each $s \in S$, $L(s) = s$, and the transition relation T contains the following transitions: $\bar{s} \longrightarrow \rho$ with $\rho = \{(r, 0.3), (y, 0.1), (g, 0.6)\}$, $r \longrightarrow \delta_{\blacksquare}$, $y \longrightarrow \delta_{\blacktriangle}$, $g \longrightarrow \delta_{\bullet}$, $r \longrightarrow \delta_{\bar{s}}$, and $g \longrightarrow \delta_{\bar{s}}$. ♦

2.2.1. Synchronous Product

The following definition of synchronous product is a variation of the definition of parallel composition provided in [20, 21], where the synchronization occurs for each pair of enabled transitions. This corresponds to the original definition of parallel composition for probabilistic automata having all transitions labelled by the same external action.

Definition 3. Given two *PAs* $\mathfrak{P}_1$ and $\mathfrak{P}_2$, the *synchronous product* of $\mathfrak{P}_1$ and $\mathfrak{P}_2$, denoted by $\mathfrak{P}_1 \otimes \mathfrak{P}_2$, is the probabilistic automaton $\mathfrak{P} = (S, \bar{s}, \mathtt{AP}, L, T)$ where $S = S_1 \times S_2$; $\bar{s} = (\bar{s}_1, \bar{s}_2)$; $\mathtt{AP} = \mathtt{AP}_1 \cup \mathtt{AP}_2$; for each $(s_1, s_2) \in S$, $L(s_1, s_2) = L_1(s_1) \cup L_2(s_2)$; and $T = \{\, ((s_1, s_2), \mu_1 \times \mu_2) \mid (s_1, \mu_1) \in T_1 \text{ and } (s_2, \mu_2) \in T_2 \,\}$.

For two *PAs* $\mathfrak{P}_1$ and $\mathfrak{P}_2$ and their synchronous product $\mathfrak{P}_1 \otimes \mathfrak{P}_2$, we refer to $\mathfrak{P}_1$ and $\mathfrak{P}_2$ as the component automata and to $\mathfrak{P}_1 \otimes \mathfrak{P}_2$ as the product automaton.

2.2.2. Probabilistic Bisimulation

As for the definition of synchronous product, the following definition of (strong) probabilistic bisimulation is a variation of the definition provided in [21], where all actions are treated as being the same external action. We first introduce the definition of combined transition.

Definition 4. Given a *PA* $\mathfrak{P}$ and a state s, we say that s enables a combined transition reaching the distribution μ, denoted by $s \longrightarrow_{\mathrm{c}} \mu$, if there exist a finite set of indexes I, a multiset of transitions $\{\, (s, \mu_i) \in T \mid i \in I \,\}$, and a multiset of real values $\{\, p_i \in \mathbb{R}_{\geq 0} \mid i \in I \,\}$ such that $\sum_{i \in I} p_i = 1$ and $\mu = \sum_{i \in I} p_i \cdot \mu_i$.

Definition 5. Given a *PA* $\mathfrak{P}$, an equivalence relation $\mathcal{R} \subseteq S \times S$ is a *(strong) (action agnostic) probabilistic bisimulation* on $\mathfrak{P}$ if, for each $(s, t) \in \mathcal{R}$, $L(s) = L(t)$ and for each $s \longrightarrow \mu_s$, there exists a combined transition $t \longrightarrow_{\mathrm{c}} \mu_t$ such that $\mu_s \; \mathcal{L}(\mathcal{R}) \; \mu_t$.

Given two states s and t, we say that s and t are probabilistically bisimilar, denoted by $s \sim^{p}_{aa} t$, if there exists a probabilistic bisimulation $\mathcal{R}$ on $\mathfrak{P}$ such that $(s, t) \in \mathcal{R}$.

Given two *PAs* $\mathfrak{P}_1$ and $\mathfrak{P}_2$, we say that $\mathfrak{P}_1$ and $\mathfrak{P}_2$ are probabilistically bisimilar, denoted by $\mathfrak{P}_1 \sim^{p}_{aa} \mathfrak{P}_2$, if there exists a probabilistic bisimulation $\mathcal{R}$ on the disjoint union of $\mathfrak{P}_1$ and $\mathfrak{P}_2$ such that $(\bar{s}_1, \bar{s}_2) \in \mathcal{R}$.

Proposition 1. *Given three PAs* $\mathfrak{P}_1$, $\mathfrak{P}_2$, *and* $\mathfrak{P}_3$, *if* $\mathfrak{P}_1 \sim^{p}_{aa} \mathfrak{P}_2$, *then* $\mathfrak{P}_1 \otimes \mathfrak{P}_3 \sim^{p}_{aa} \mathfrak{P}_2 \otimes \mathfrak{P}_3$.

Proof. The proof is a minor adaptation of the corresponding proof (cf. [20]) for the original definition of probabilistic bisimulation and parallel composition of *PAs*.

In the following, we use the subscript "$j, 3$" with $j \in \{1, 2\}$ to refer to the component of the *PA* $\mathfrak{P}_{j,3} = \mathfrak{P}_j \otimes \mathfrak{P}_3$.

Let $\mathcal{R}$ be the probabilistic bisimulation justifying $\mathfrak{P}_1 \sim^{p}_{aa} \mathfrak{P}_2$ and $\mathcal{R}' = \mathcal{R} \times \mathcal{I}_{S_3}$; we claim that $\mathcal{R}'$ is a probabilistic bisimulation between $\mathfrak{P}_1 \otimes \mathfrak{P}_3$ and $\mathfrak{P}_2 \otimes \mathfrak{P}_3$. The fact that $\mathcal{R}'$ is an equivalence relation follows trivially by its definition and the fact that $\mathcal{R}$ is an equivalence relation. The fact that $((\bar{s}_1, \bar{s}_3), (\bar{s}_2, \bar{s}_3))$ follows immediately by the hypothesis that $(\bar{s}_1, \bar{s}_2) \in \mathcal{R}$ and $(\bar{s}_3, \bar{s}_3) \in \mathcal{I}_{S_3}$.

Let $((s_1, s_3), (s_2, s_3)) \in \mathcal{R}'$. Assume, without loss of generality, that $s_1 \in S_1$ and $s_2 \in S_2$; the other cases are similar. The fact that $L_{1,3}(s_1, s_3) = L_{2,3}(s_2, s_3)$ is straightforward, since by definition of synchronous product and the hypothesis that $s_1 \; \mathcal{R} \; s_2$, we have that $L_{1,3}(s_1, s_3) = L_1(s_1) \cup L_3(s_3) = L_2(s_2) \cup L_3(s_3) = L_{2,3}(s_2, s_3)$, as required.

Consider now a transition $(s_1, s_3) \longrightarrow \mu_{1,3}$. By definition of synchronous product, there exist μ_1 and μ_3 such that $s_1 \longrightarrow \mu_1 \in T_1$, $s_3 \longrightarrow \mu_3 \in T_3$, and $\mu_{1,3} = \mu_1 \times \mu_3$. Since $s_1 \; \mathcal{R} \; s_2$, it follows that there exists a combined transition $s_2 \longrightarrow_{\mathrm{c}} \mu_2$ such that $\mu_1 \; \mathcal{L}(\mathcal{R}) \; \mu_2$. Let I be a the finite set of indexes, $\{\, (s_2, \mu_{2,i}) \in T_2 \mid i \in I \,\}$ be a multiset of transitions, and $\{\, p_i \in \mathbb{R}_{\geq 0} \mid i \in I \,\}$ be a multiset of real values such that $\sum_{i \in I} p_i = 1$ and $\mu_2 = \sum_{i \in I} p_i \cdot \mu_{2,i}$. By definition of synchronous product, it follows that for each $i \in I$, $(s_2, s_3) \longrightarrow \mu_{2,i} \times \mu_3 \in T_{2,3}$, hence we have the combined transition $(s_2, s_3) \longrightarrow_{\mathrm{c}} \mu_2 \times \mu_3$. By standard properties of lifting (see, e.g., [23]), it follows that $\mu_1 \times \mu_3 \; \mathcal{L}(\mathcal{R}') \; \mu_2 \times \mu_3$, as required. □

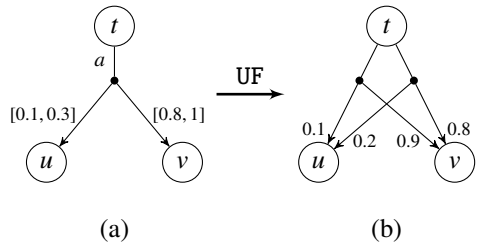

Figure 2: Unfolding *IMDP* $\mathfrak{I}$ to *PA* $\mathfrak{P}$

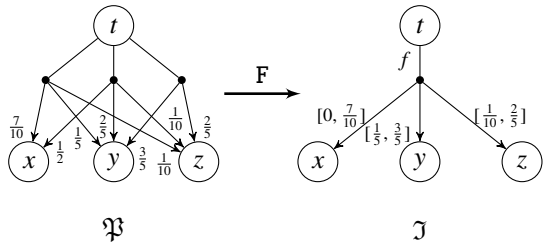

Figure 3: Folding a *PA* $\mathfrak{P}$ as an *IMDP* $\mathfrak{I}$

2.3. IMDPs vs. PAs

A cornerstone towards establishing compositional reasoning for *IMDPs* essentially relies on *transformations* from *IMDPs* to *PAs* and vice versa. To this aim, we define two mappings namely, *unfolding* which unfolds a given *IMDP* as a *PA* and *folding* which transforms a given *PA* to an *IMDP*. Formally,

Definition 6 (Unfolding mapping). An unfolding mapping $\mathtt{UF}\colon [\mathfrak{I}] \to [\mathfrak{P}]$ is a function that maps a given *IMDP* $\mathfrak{I} = (S, \bar{s}, \mathcal{A}, \mathtt{AP}, L, I)$ to the *PA* $\mathfrak{P} = (S, \bar{s}, \mathtt{AP}, L, T)$ where $T = \{(s,\mu) \mid s \in S, \exists a \in \mathcal{A}(s) : \mu \in \mathtt{Ext}(\mathcal{P}^{s,a})\}$.

It is worthy to note that the unfolding mapping might transform an *IMDP* to a *PA* with an exponentially larger size. This is in fact due to the exponential blow up in the number of transitions in the resultant *PA* which in turn depends on the number of extreme points of the polytope constructed for each state and action in the given *IMDP*. An example of unfolding is given in Figure 2.

In order to transform a given *PA* to an instance of *IMDPs*, we use the folding mapping defined as follows:

Definition 7 (Folding mapping). The folding mapping $\mathtt{F}\colon [\mathfrak{P}] \to [\mathfrak{I}]$ transforms a *PA* $\mathfrak{P} = (S, \bar{s}, \mathtt{AP}, L, T)$ to the *IMDP* $\mathfrak{I} = (S, \bar{s}, \{f\}, \mathtt{AP}, L, I)$ where, for each $s, t \in S$, $I(s, f, t) = \mathrm{proj}_{\mathbf{e}_t}\,\mathrm{CH}(\{\mu \mid (s,\mu) \in T\})$, where each component $\mathbf{e}_{\mathbf{uv}}$ of the vector $\mathbf{e}_{\mathbf{u}} \in \mathbb{R}^{|S|}$ is defined as $\mathbf{e}_{\mathbf{uv}} = \delta_u(v)$.

An example of the folding mapping is shown in Figure 3. The *PA* $\mathfrak{P}$ has three transitions from t with label a; in particular, it is worthwhile to note that for all these transitions the probability of reaching y is larger than the probability of reaching z, so this has to happen for every combined transition leaving t. According to Def. 7, the folding of $\mathfrak{P}$ is the *IMDP* $\mathfrak{I}$. It is immediate to see that the unfolding mapping is not surjective as there may be some probabilistic transitions in the generated

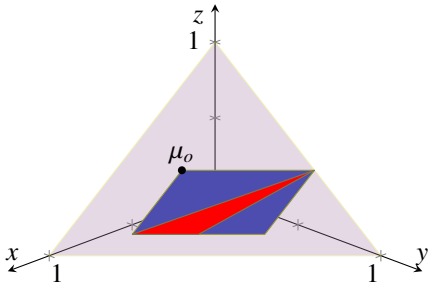

Figure 4: Comparison of polytopes resulted from folding mapping F

IMDP specification which cannot be mapped to a probability distribution in the given *PA*. In fact, one of such distributions is μ_o such that $\mu_o(x) = \frac{2}{5}$, $\mu_o(y) = \frac{1}{5}$, and $\mu_o(z) = \frac{2}{5}$ that clearly violates the condition $\mu_o(y) > \mu_o(z)$. This is better recognizable by comparing the corresponding polytopes in a graphical way. Figure 4 shows the three polytopes involved in $\mathfrak{I}$: the purplish large triangular polytope is the standard *2-simplex* in the three dimensional space; the reddish small triangular and the bluish parallelogram-like polytopes represent the convex hull of $\left\{\left(\frac{7}{10}, \frac{1}{5}, \frac{1}{10}\right), \left(\frac{1}{2}, \frac{2}{5}, \frac{1}{10}\right), \left(0, \frac{3}{5}, \frac{2}{5}\right)\right\}$ and the polytope $\mathcal{P}^{t,f}$, respectively, both being a sub-polytope of the 2-simplex. Clearly there are points in $\mathcal{P}^{t,f}$ that do not belong to the reddish polytope, such as the black dot corresponding to μ_o.

Lemma 1. *Given the folding and unfolding mappings* $\mathtt{F}\colon [\mathfrak{P}] \to [\mathfrak{I}]$ *and* $\mathtt{UF}\colon [\mathfrak{I}] \to [\mathfrak{P}]$*, in general:*

1. $\mathtt{UF}(\mathtt{F}(\mathfrak{P})) \neq \mathfrak{P}$
2. $\mathtt{F}(\mathtt{UF}(\mathfrak{I})) \neq \mathfrak{I}$

As we will discuss later, the general incompleteness property of the folding mapping does not influence on the generality of our compositional reasoning for *IMDP* specifications. We will dive into this point later in Section 4.

3. Probabilistic Bisimulation for Interval MDPs

We now recall the main results on probabilistic bisimulation for *IMDPs*, as developed in [8]. In this work, we consider the notion of probabilistic bisimulation for the cooperative resolution of nondeterminism. This semantics is very natural in the context of verification of parallel systems with uncertain transition probabilities in which we assume that scheduler and nature are resolved *cooperatively* in the most *adversarial* way. Moreover, resolution of a feasible probability distribution respecting the interval constraints can be either done *statically* [13], i.e., at the beginning once for all, or *dynamically* [12, 22], i.e., independently for each computation step. In this paper, we focus on dynamic approach in resolving the stochastic nondeterminism that is easier to work with algorithmically and can be seen as a relaxation of the static approach that is often intractable [1, 3].

Let $s \longrightarrow \mu_s$ denote that a transition from s to μ_s can be taken cooperatively, i.e., that there is a scheduler $\sigma \in \Sigma$ and a nature

$\pi \in \Pi$ such that $\mu_s = \sum_{a \in \mathcal{A}(s)} \sigma(s)(a) \cdot \pi(s,a)$. In other words, $s \longrightarrow \mu_s$ if $\mu_s \in \mathrm{CH}(\bigcup_{a \in \mathcal{A}(s)} \mathcal{P}^{s,a})$.

Definition 8 (cf. [8]). Given an *IMDP* $\mathfrak{I}$, let $\mathcal{R} \subseteq S \times S$ be an equivalence relation. We say that $\mathcal{R}$ is a *probabilistic bisimulation* if for each $(s,t) \in \mathcal{R}$ we have that $L(s) = L(t)$ and for each $s \longrightarrow \mu_s$ there exists $t \longrightarrow \mu_t$ such that $\mu_s \ \mathcal{L}(\mathcal{R}) \ \mu_t$. Furthermore, we write $s \sim_c t$ if there is a probabilistic bisimulation $\mathcal{R}$ such that $(s,t) \in \mathcal{R}$.

Intuitively, each (cooperative) step of scheduler and nature from state s needs to be matched by a (cooperative) step of scheduler and nature from state t; symmetrically, s also needs to match t. In order to support the compositional reasoning, $\sim_c$ needs to be an equivalence relation. It is not difficult to see that $\sim_c$ is reflexive and symmetric. What remains is to show that it is also transitive. This is indeed a property of $\sim_c$, as stated by the following proposition:

Theorem 1. *Given three IMDPs* $\mathfrak{I}_1$, $\mathfrak{I}_2$, *and* $\mathfrak{I}_3$, *if* $\mathfrak{I}_1 \sim_c \mathfrak{I}_2$ *and* $\mathfrak{I}_2 \sim_c \mathfrak{I}_3$, *then* $\mathfrak{I}_1 \sim_c \mathfrak{I}_3$.

Proof. Let $\mathcal{R}_{12}$ and $\mathcal{R}_{23}$ be the equivalence relations underlying $\mathfrak{I}_1 \sim_c \mathfrak{I}_2$ and $\mathfrak{I}_2 \sim_c \mathfrak{I}_3$, respectively. Let $\mathcal{R}_{13}$ be the symmetric and transitive closure of the set $\{(s_1,s_3) \mid \exists s_2 . s_1 \ \mathcal{R}_{12} \ s_2 \wedge s_2 \ \mathcal{R}_{23} \ s_3\} \cup \{(s_1,s_1') \in S_1 \times S_1 \mid (s_1,s_1') \in \mathcal{R}_{12}\} \cup \{(s_3,s_3') \in S_3 \times S_3 \mid (s_3,s_3') \in \mathcal{R}_{23}\}$. We claim that $\mathcal{R}_{13}$ is a probabilistic bisimulation justifying $\mathfrak{I}_1 \sim_c \mathfrak{I}_3$.

The fact that $\bar{s}_1 \ \mathcal{R}_{13} \ \bar{s}_3$ is trivial since by hypothesis we have that $\bar{s}_1 \ \mathcal{R}_{12} \ \bar{s}_2$ and $\bar{s}_2 \ \mathcal{R}_{23} \ \bar{s}_3$, so $(\bar{s}_1, \bar{s}_3) \in \mathcal{R}_{13}$ by construction.

In the following, assume that $s_1 \in S_1$ and $s_3 \in S_3$; the other cases are similar.

The labelling is respected: for each $s_1 \ \mathcal{R}_{13} \ s_3$, we have that there exists s_2 such that $s_1 \ \mathcal{R}_{12} \ s_2$ and $s_2 \ \mathcal{R}_{23} \ s_3$; this implies that $L_1(s_1) = L_2(s_2)$ and $L_2(s_2) = L_3(s_3)$, thus $L_1(s_1) = L_3(s_3)$ as required.

To complete the proof, consider $s_1 \ \mathcal{R}_{13} \ s_3$ and $s_1 \longrightarrow \mu_1$. By hypothesis, there exists s_2 such that $s_1 \ \mathcal{R}_{12} \ s_2$ and $s_2 \ \mathcal{R}_{23} \ s_3$; moreover, by $\mathcal{R}_{12}$ being a probabilistic bisimulation, we know that there exists $s_2 \longrightarrow \mu_2$ such that $\mu_1 \ \mathcal{L}(\mathcal{R}_{12}) \ \mu_2$. Since $\mathcal{R}_{23}$ is a probabilistic bisimulation, we have that there exists $s_3 \longrightarrow \mu_3$ such that $\mu_2 \ \mathcal{L}(\mathcal{R}_{23}) \ \mu_3$. By construction of $\mathcal{R}_{13}$ and the properties of lifting, it follows that $\mu_1 \ \mathcal{L}(\mathcal{R}_{13}) \ \mu_3$, as required. □

It is shown in [8] that $\sim_c$ is sound with respect to the PCTL properties. Furthermore, probabilistic bisimulation for *IMDPs* is computed using standard partition refinement approach [14, 19] in which the core part is to verify the violation of bisimulation definition that can in turn be done by checking the inclusion of polytopes defined as follows. For $s \in S$ and an action $a \in \mathcal{A}$, recall that $\mathcal{P}^{s,a}$ denotes the polytope of feasible successor distributions over *states* with respect to taking the action a in the state s. By $\mathcal{P}^{s,a}_{\mathcal{R}}$, we denote the polytope of feasible successor distributions over *equivalence classes* of $\mathcal{R}$ with respect to taking the action a in the state s. Formally, for $\mu \in \Delta(S/\mathcal{R})$ we set $\mu \in \mathcal{P}^{s,a}_{\mathcal{R}}$ if, for each $C \in S/\mathcal{R}$, it is

$$\mu(C) \in \left[\sum_{s' \in C} \inf I(s,a,s'), \sum_{s' \in C} \sup I(s,a,s') \right].$$

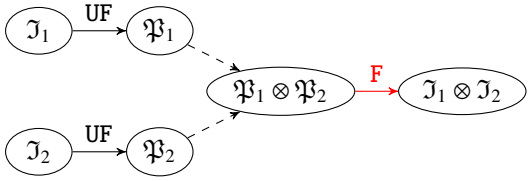

Figure 5: Schematic representation of the synchronous product of *IMDPs* $\mathfrak{I}_1$ and $\mathfrak{I}_2$. In this figure, we let $\mathfrak{P}_1 = \mathtt{UF}(\mathfrak{I}_1)$ and $\mathfrak{P}_2 = \mathtt{UF}(\mathfrak{I}_2)$

Furthermore, we define $\mathcal{P}^{s}_{\mathcal{R}} = \mathrm{CH}(\bigcup_{a \in \mathcal{A}(s)} \mathcal{P}^{s,a}_{\mathcal{R}})$, the set of feasible successor distributions over $S/\mathcal{R}$ with respect to taking an *arbitrary* distribution over enabled actions in state s. As specified in [8], checking violation of a given pair of states amounts to check equality of the corresponding constructed polytopes for the states.

4. Compositional Reasoning for *IMDPs*

The compositional reasoning is a widely used technique (see, e.g., [4, 11, 15]) that permits to deal with large systems. In particular, a large system is decomposed into multiple components running in parallel; such components are then minimized by replacing each of them by a bisimilar but smaller one so that the overall behaviour remains unchanged. In order to apply this technique, bisimulation has first to be extended to pairs of components and then to be shown to be transitive and preserved by the synchronous product operator. The extension to a pair of components is trivial and commonly done (see, e.g., [2, 20]):

Definition 9. Given two *IMDPs* $\mathfrak{I}_1$ and $\mathfrak{I}_2$, we say that they are *probabilistic bisimilar*, denoted by $\mathfrak{I}_1 \sim_c \mathfrak{I}_2$, if there exists a probabilistic bisimulation on the disjoint union of $\mathfrak{I}_1$ and $\mathfrak{I}_2$ such that $\bar{s}_1 \sim_c \bar{s}_2$.

The next step is to define the synchronous product for *IMDPs*:

Definition 10. Given two *IMDPs* $\mathfrak{I}_1$ and $\mathfrak{I}_2$, we define the *synchronous product* of $\mathfrak{I}_1$ and $\mathfrak{I}_2$ as

$$\mathfrak{I}_1 \otimes \mathfrak{I}_2 := \mathtt{F}(\mathtt{UF}(\mathfrak{I}_1) \otimes \mathtt{UF}(\mathfrak{I}_2)).$$

A schematic representation of constructing the synchronous product of two *IMDPs* $\mathfrak{I}_1$ and $\mathfrak{I}_2$ is given in Figure 5. As discussed earlier, the folding mapping from *PA* to *IMDP*, i.e. the red arrow, is not complete and in principle, this transformation may add additional behavior to the resultant system. For each state and action in the resultant *IMDP*, these extra behaviors are essentially a set of probability distributions that do not belong to the convex hull of the enabled probability distributions for that state in the original *PA*. At first sight, these extra behaviors generated from the folding mapping might be seen as an impediment towards showing that $\sim_c$ is a congruence for the synchronous product. Fortunately, as it is shown by the next theorem, these extra probability distributions are in fact *spurious* and do not affect the congruence result. To this aim and in order to pave the way for establishing the congruence result, we

first prove two intermediate results stating that the folding and unfolding mappings preserve bisimilarity on the corresponding codomains.

Lemma 2. *Given two IMDPs* $\mathfrak{I}_1$ *and* $\mathfrak{I}_2$*, if* $\mathfrak{I}_1 \sim_c \mathfrak{I}_2$*, then* $\mathtt{UF}(\mathfrak{I}_1) \sim^p_{aa} \mathtt{UF}(\mathfrak{I}_2)$.

Proof. Let $\mathcal{R}$ be the probabilistic bisimulation justifying $\mathfrak{I}_1 \sim_c \mathfrak{I}_2$; we claim that $\mathcal{R}$ is also a *PA* probabilistic bisimulation for $\mathtt{UF}(\mathfrak{I}_1)$ and $\mathtt{UF}(\mathfrak{I}_2)$, that is, it justifies $\mathtt{UF}(\mathfrak{I}_1) \sim^p_{aa} \mathtt{UF}(\mathfrak{I}_2)$.

In the following we assume without loss of generality that $s_1 \in S_1$ and $s_2 \in S_2$; the other cases are similar. The fact that $\mathcal{R}$ is an equivalence relation and that for each $(s_1, s_2) \in \mathcal{R}$, $L_1(s_1) = L_2(s_2)$ follow directly by definition of $\sim_c$. Let $(s_1, \mu_1) \in T_1$: by definition of $\mathtt{UF}$, it follows that $\mu_1 \in \mathtt{Ext}(\mathcal{P}^{s_1,a_1})$ for some $a_1 \in \mathcal{A}(s_1)$, thus in particular $\mu_1 \in \mathcal{P}^{s_1,a_1}$, hence $\mu_1 \in \mathrm{CH}(\bigcup_{a \in \mathcal{A}(s_1)} \mathcal{P}^{s_1,a})$. By hypothesis, we have that there exists $\mu_2 \in \mathrm{CH}(\bigcup_{a_2 \in \mathcal{A}(s_2)} \mathcal{P}^{s_2,a_2})$ such that $\mu_1 \; \mathcal{L}(\mathcal{R}) \; \mu_2$. Since $\mu_2 \in \mathrm{CH}(\cup_{a_2 \in \mathcal{A}(s_2)} \mathcal{P}^{s_2,a_2})$, it follows that there exist a multiset of real values $\{\, p_{a_2} \in \mathbb{R}_{\geq 0} \mid a_2 \in \mathcal{A}(s_2) \,\}$ and a multiset of distributions $\{\, \mu_{a_2} \in \mathcal{P}^{s_2,a_2} \mid a_2 \in \mathcal{A}(s_2) \,\}$ such that $\sum_{a_2 \in \mathcal{A}(s_2)} p_{a_2} = 1$ and $\mu_2 = \sum_{a_2 \in \mathcal{A}(s_2)} p_{a_2} \cdot \mu_{a_2}$. For each $a_2 \in \mathcal{A}(s_2)$, since $\mu_{a_2} \in \mathcal{P}^{s_2,a_2}$, it follows that there exist a finite set of indexes I_{a_2}, a multiset of real values $\{\, p_{a_2,i} \in \mathbb{R}_{\geq 0} \mid i \in I_{a_2} \,\}$ and a multiset of distributions $\{\, \mu_{a_2,i} \in \mathtt{Ext}(\mathcal{P}^{s_2,a_2}) \mid i \in I_{a_2} \,\}$ such that $\sum_{i \in I_{a_2}} p_{a_2,i} = 1$ and $\mu_{a_2} = \sum_{i \in I_{a_2}} p_{a_2,i} \cdot \mu_{a_2,i}$. This means that $\mu_2 = \sum_{a_2 \in \mathcal{A}(s_2)} p_{a_2} \cdot \sum_{i \in I_{a_2}} p_{a_2,i} \cdot \mu_{a_2,i} = \sum_{a_2 \in \mathcal{A}(s_2)} \sum_{i \in I_{a_2}} p_{a_2} \cdot p_{a,i} \cdot \mu_{a_2,i}$. Since for each $a_2 \in \mathcal{A}(s_2)$ and $i \in I_{a_2}$ we have that $\mu_{a_2,i} \in \mathtt{Ext}(\mathcal{P}^{s_2,a_2})$, it follows that $(s_2, \mu_{a_2,i}) \in T_2$, thus we have the combined transition $s_2 \longrightarrow_{\mathrm{c}} \mu_2$ obtained by taking as set of indexes $I = \{\, (a_2, i) \mid a_2 \in \mathcal{A}(s_2), i \in I_{a_2} \,\}$, as multiset of real values $\{\, q_{a_2,i} \in \mathbb{R}_{\geq 0} \mid (a_2, i) \in I, q_{a_2,i} = p_{a_2} \cdot p_{a_2,i} \,\}$, and as multiset of transitions $\{\, (s_2, \mu_{a_2,i}) \in T_2 \mid (a_2, i) \in I \,\}$: in fact, it is immediate to see that

$$\sum_{(a_2,i) \in I} q_{a_2,i} = \sum_{(a_2,i) \in I} p_{a_2} \cdot p_{a_2,i} = \sum_{a_2 \in \mathcal{A}(s_2)} \sum_{i \in I_{a_2}} p_{a_2} \cdot p_{a_2,i}$$
$$= \sum_{a_2 \in \mathcal{A}(s_2)} p_{a_2} \cdot \sum_{i \in I_{a_2}} p_{a_2,i} = \sum_{a_2 \in \mathcal{A}(s_2)} p_{a_2} \cdot 1 = 1$$

and that

$$\sum_{(a_2,i) \in I} q_{a_2,i} \cdot \mu_{a_2,i} = \sum_{(a_2,i) \in I} p_{a_2} \cdot p_{a_2,i} \cdot \mu_{a_2,i}$$
$$= \sum_{a_2 \in \mathcal{A}(s_2)} \sum_{i \in I_{a_2}} p_{a_2} \cdot p_{a_2,i} \cdot \mu_{a_2,i} = \sum_{a_2 \in \mathcal{A}(s_2)} p_{a_2} \cdot \sum_{i \in I_{a_2}} p_{a_2,i} \cdot \mu_{a_2,i}$$
$$= \sum_{a_2 \in \mathcal{A}(s_2)} p_{a_2} \cdot \mu_{a_2} = \mu_2.$$

Moreover, by hypothesis, we have $\mu_1 \; \mathcal{L}(\mathcal{R}) \; \mu_2$, as required. □

Likewise computation of probabilistic bisimulation for *IMDPs*, we use the standard *partition refinement* approach as a ground procedure to compute $\sim^p_{aa}$ for *PAs*. Still the core part of the approach is to decide bisimilarity of a pair of states. For each state in the given *PA*, we construct a convex hull polytope which encodes all possible behaviors that can be taken by a scheduler. Hence, for a given pair of states, we show that verifying if two states are bisimilar can be reduced to comparison of their corresponding convex polytopes with respect to set inclusion. Strictly speaking, for an equivalence relation $\mathcal{R}$ on S and $s \in S$, we denote by $\mathrm{P}^s_{\mathcal{R}}$ the polytope of feasible successor distributions over *equivalence classes* of $\mathcal{R}$ with respect to taking a transition in the state s. Formally,

$$\mathrm{P}^s_{\mathcal{R}} = \mathrm{CH}(\{\, [\mu]_{\mathcal{R}} \mid (s, \mu) \in T \,\}),$$

where, for a given $\mu \in \Delta(S)$, $[\mu]_{\mathcal{R}} \in \Delta(S/\mathcal{R})$ is the probability distribution such that for each $C \in S/\mathcal{R}$, it is $[\mu]_{\mathcal{R}}(C) = \sum_{s' \in C} \mu(s')$.

Lemma 3 (cf. [2, Thm. 1])**.** *Given a PA* $\mathfrak{P}$*, there exists an equivalence relation* $\mathcal{R}$ *on* S *such that for each pair states* $s, t \in S$*, it holds that* $s \sim^p_{aa} t$ *if and only if* $s \; \mathcal{R} \; t$*,* $L(s) = L(t)$*, and* $\mathrm{P}^s_{\mathcal{R}} = \mathrm{P}^t_{\mathcal{R}}$.

To simplify the presentation of the proof, we first introduce some notation. Given an equivalence relation $\mathcal{R}$ on S, for each distribution $\mu \in \Delta(S)$, let $\bar{\mu} \in \Delta(S/\mathcal{R})$ denote the corresponding distribution $\bar{\mu} = [\mu]_{\mathcal{R}}$, i.e., $\bar{\mu}$ is such that $\bar{\mu}(C) = \sum_{s' \in C} \mu(s')$ for each $C \in S/\mathcal{R}$.

Proof. We show the two implications separately. For the implication from left to right, suppose that $s \sim^p_{aa} t$; this implies that there exists a probabilistic bisimulation $\mathcal{R}$ such that $s \; \mathcal{R} \; t$ and $L(s) = L(t)$. We want to show that $\mathrm{P}^s_{\mathcal{R}} = \mathrm{P}^t_{\mathcal{R}}$ holds. To this aim, let $\eta \in \mathrm{P}^s_{\mathcal{R}}$. By definition of $\mathrm{P}^s_{\mathcal{R}}$, it follows that there exist a finite set of indexes I_η, a multiset of real values $\{\, p_{\eta,i} \mid i \in I_\eta \,\}$ and a multiset of distributions $\{\, \eta_i \in \mathrm{P}^s_{\mathcal{R}} \mid i \in I_\eta, \exists (s, \mu_{s,i}) \in T : \eta_i = [\mu_{s,i}]_{\mathcal{R}} \,\}$ such that $\sum_{i \in I_\eta} p_{\eta,i} = 1$ and $\sum_{i \in I_\eta} p_{\eta,i} \cdot \eta_i = \eta$. Since $s \; \mathcal{R} \; t$ and $\mathcal{R}$ is a probabilistic bisimulation, it follows that for each $i \in I_\eta$ there exists a combined transition $t \longrightarrow_{\mathrm{c}} \mu_t$ such that $\mu_{s,i} \; \mathcal{L}(\mathcal{R}) \; \mu_t$. By definition of combined transition, it follows that there exist a finite set of indexes I_t, a set of transitions $\{\, (t, \mu_{t,i}) \in T \mid i \in I_t \,\}$ and a multiset of real values $\{\, p_{t,i} \in \mathbb{R}_{\geq 0} \mid i \in I_t \,\}$ such that $\sum_{i \in I_t} p_{t,i} = 1$ and $\mu_t = \sum_{i \in I_t} p_{t,i} \cdot \mu_{t,i}$. This implies that for each $i \in I_t$, $\bar{\mu}_{t,i} \in \mathrm{P}^t_{\mathcal{R}}$. Moreover, since by definition of lifting we have that for each $C \in S/\mathcal{R}$, $\mu_s(C) = \mu_t(C)$, it follows immediately that $\bar{\mu}_t = \bar{\mu}_s$, thus we have that $\eta = \bar{\mu}_s = \bar{\mu}_t \in \mathrm{P}^t_{\mathcal{R}}$, hence $\mathrm{P}^s_{\mathcal{R}} \subseteq \mathrm{P}^t_{\mathcal{R}}$. By swapping the roles of s and t, we can show in the same way that $\mathrm{P}^t_{\mathcal{R}} \subseteq \mathrm{P}^s_{\mathcal{R}}$, hence $\mathrm{P}^s_{\mathcal{R}} = \mathrm{P}^t_{\mathcal{R}}$ as required.

For the implication from right to left, fix an equivalence relation $\mathcal{R}$ on S such that for each $(s, t) \in \mathcal{R}$ it holds that $L(s) = L(t)$ and $\mathrm{P}^s_{\mathcal{R}} = \mathrm{P}^t_{\mathcal{R}}$; we want to show that $\mathcal{R}$ is a probabilistic bisimulation, i.e., whenever $s \; \mathcal{R} \; t$ and $s \longrightarrow \mu_s$ then there exists $t \longrightarrow_{\mathrm{c}} \mu_t$ such that $\mu_s \; \mathcal{L}(\mathcal{R}) \; \mu_t$. Let $(s, t) \in \mathcal{R}$; if $\mathrm{P}^s_{\mathcal{R}} = \emptyset$, then the step condition of the probabilistic bisimulation is trivially verified since there is no transition $s \longrightarrow \mu_s$ from s that needs to be matched by t. Suppose now that $\mathrm{P}^s_{\mathcal{R}} \neq \emptyset$ and consider a transition $s \longrightarrow \mu_s$ so that $\bar{\mu}_s \in \mathrm{P}^s_{\mathcal{R}}$. By hypothesis, $\bar{\mu}_s \in \mathrm{P}^s_{\mathcal{R}} = \mathrm{P}^t_{\mathcal{R}}$, thus there exist a finite set of indexes I, a multiset of distributions $\{\, \mu_i \in \mathrm{P}^t_{\mathcal{R}} \mid i \in I \,\}$ and a multiset of real values $\{\, p_i \in \mathbb{R}_{\geq 0} \mid i \in I \,\}$ such that $\sum_{i \in I} p_i = 1$ and $\sum_{i \in I} p_i \cdot \mu_i = \bar{\mu}_s$. This implies, for each $i \in I$, that there exist a finite set of indexes J_i, a multiset of

real values $\{ p_{i,j} \in \mathbb{R}_{\geq 0} \mid j \in I_i \}$, and a multiset of distributions $\{ \mu_{i,j} \in \mathrm{P}^t_{\mathcal{R}} \mid j \in J_i \}$ such that $\sum_{j \in J_i} p_{i,j} = 1$, $\sum_{j \in J_i} p_{i,j} \cdot \mu_{i,j} = \mu_i$, and for each $j \in J_i$, $\mu_{i,j} = \bar{\mu}_{t,i,j}$ where $(t, \mu_{t,i,j}) \in T$. Consider now the combined transition $t \longrightarrow_{\mathrm{c}} \mu_t$ obtained by taking as set of indexes $J = \{ (i,j) \mid i \in I, j \in J_i \}$, as multiset of real values $\{ q_{i,j} \in \mathbb{R}_{\geq 0} \mid (i,j) \in J, q_{i,j} = p_i \cdot p_{i,j} \}$, and as set of transitions $\{ (t, \mu_{t,i,j}) \in T \mid (i,j) \in J \}$: we have that

$$\sum_{(i,j) \in J} q_{i,j} = \sum_{(i,j) \in J} p_i \cdot p_{i,j} = \sum_{i \in I} \sum_{j \in J_i} p_i \cdot p_{i,j}$$
$$= \sum_{i \in I} p_i \cdot \sum_{j \in J_i} p_{i,j} = \sum_{i \in I} p_i \cdot 1 = 1$$

and that

$$\mu_t = \sum_{(i,j) \in I} q_{i,j} \cdot \mu_{t,i,j} = \sum_{(i,j) \in I} p_i \cdot p_{i,j} \cdot \mu_{t,i,j}$$
$$= \sum_{i \in I} \sum_{j \in J_i} p_i \cdot p_{i,j} \cdot \mu_{t,i,j} = \sum_{i \in I} p_i \cdot \sum_{j \in J_i} p_{i,j} \cdot \mu_{t,i,j}.$$

To complete the proof, we have to show that $\mu_s \; \mathcal{L}(\mathcal{R}) \; \mu_t$, that is, for each $C \in S/\mathcal{R}$, $\mu_s(C) = \mu_t(C)$. Let $C \in S/\mathcal{R}$: we have that

$$\mu_t(C) = \sum_{c \in C} \mu_t(c) = \sum_{c \in C} \sum_{i \in I} p_i \cdot \sum_{j \in J_i} p_{i,j} \cdot \mu_{t,i,j}(c)$$
$$= \sum_{i \in I} p_i \cdot \sum_{j \in J_i} p_{i,j} \cdot \sum_{c \in C} \mu_{t,i,j}(c) = \sum_{i \in I} p_i \cdot \sum_{j \in J_i} p_{i,j} \cdot \bar{\mu}_{t,i,j}(C)$$
$$= \sum_{i \in I} p_i \cdot \sum_{j \in J_i} p_{i,j} \cdot \mu_{i,j}(C) = \sum_{i \in I} p_i \cdot \mu_i(C) = \bar{\mu}_s(C)$$
$$= \sum_{c \in C} \mu_s(c) = \mu_s(C),$$

as required. □

Lemma 4. *Given a PA $\mathfrak{P}$ and an equivalence relation $\mathcal{R}$ on S, for $n = |S/\mathcal{R}|$, it holds that for each $(s,t) \in \mathcal{R}$, if $\mathrm{P}^s_{\mathcal{R}} = \mathrm{P}^t_{\mathcal{R}}$ then $(\prod_{C \in S/\mathcal{R}} \mathrm{proj}_{\mathbf{e}_C} \mathrm{P}^s_{\mathcal{R}}) \cap \Delta_n = (\prod_{C \in S/\mathcal{R}} \mathrm{proj}_{\mathbf{e}_C} \mathrm{P}^t_{\mathcal{R}}) \cap \Delta_n$.*

Proof. The proof is trivial, since by $\mathrm{P}^s_{\mathcal{R}} = \mathrm{P}^t_{\mathcal{R}}$ it follows that for each $C \in S/\mathcal{R}$, $\mathrm{proj}_{\mathbf{e}_C} \mathrm{P}^s_{\mathcal{R}} = \mathrm{proj}_{\mathbf{e}_C} \mathrm{P}^t_{\mathcal{R}}$. This implies that $\prod_{C \in S/\mathcal{R}} \mathrm{proj}_{\mathbf{e}_C} \mathrm{P}^s_{\mathcal{R}} = \prod_{C \in S/\mathcal{R}} \mathrm{proj}_{\mathbf{e}_C} \mathrm{P}^t_{\mathcal{R}}$ thus $(\prod_{C \in S/\mathcal{R}} \mathrm{proj}_{\mathbf{e}_C} \mathrm{P}^s_{\mathcal{R}}) \cap \Delta_n = (\prod_{C \in S/\mathcal{R}} \mathrm{proj}_{\mathbf{e}_C} \mathrm{P}^t_{\mathcal{R}}) \cap \Delta_n$, as required. □

Lemma 5. *Given two PAs $\mathfrak{P}_1$ and $\mathfrak{P}_2$, if $\mathfrak{P}_1 \sim^p_{aa} \mathfrak{P}_2$ then $\mathtt{F}(\mathfrak{P}_1) \sim_c \mathtt{F}(\mathfrak{P}_2)$.*

Proof. Let $\mathcal{R}$ be the equivalence relation justifying $\mathfrak{P}_1 \sim^p_{aa} \mathfrak{P}_2$; we claim that $\mathcal{R}$ is also an *IMDP* probabilistic bisimulation for $\mathtt{F}(\mathfrak{P}_1)$ and $\mathtt{F}(\mathfrak{P}_2)$, that is, it justifies $\mathtt{F}(\mathfrak{P}_1) \sim_c \mathtt{F}(\mathfrak{P}_2)$.

In the following we assume without loss of generality that $s_1 \in S_1$ and $s_2 \in S_2$; the other cases are similar. The fact that $\mathcal{R}$ is an equivalence relation and that for each $(s_1, s_2) \in \mathcal{R}$, $L_1(s_1) = L_2(s_2)$ follow directly by definition of $\sim^p_{aa}$. Since $\mathfrak{P}_1 \sim^p_{aa} \mathfrak{P}_2$, it follows from Lemma 3 that $\mathrm{P}^{s_1}_{\mathcal{R}} = \mathrm{P}^{s_2}_{\mathcal{R}}$. Additionally, it is not difficult to see that for $s_j \in \{s_1, s_2\}$, $\mathcal{P}^{s_j}_{\mathcal{R}} = \mathcal{P}^{s_j,f}_{\mathcal{R}} = (\prod_{i=1}^n \mathrm{proj}_{\mathbf{e}_i} \mathrm{P}^{s_j}_{\mathcal{R}}) \cap \Delta_n$ where $\Delta_n = \{ (x_1, \ldots, x_n) \in \mathbb{R}^n_{\geq 0} \mid \sum_{i=1}^n x_i = 1 \}$ is the standard probability simplex in $\mathbb{R}^n$. By Lemma 4, this implies that $\mathcal{P}^{s_1,f}_{\mathcal{R}} = \mathcal{P}^{s_2,f}_{\mathcal{R}}$. Consider now $s_1 \longrightarrow \mu_1$ with $\mu_1 \in \mathcal{P}^{s_1,f}$: this implies that $[\mu_1]_{\mathcal{R}} \in \mathcal{P}^{s_1,f}_{\mathcal{R}}$ since $\mathcal{P}^{s_1,f}_{\mathcal{R}} = \mathcal{P}^{s_2,f}_{\mathcal{R}}$, it follows that $[\mu_1]_{\mathcal{R}} \in \mathcal{P}^{s_2,f}_{\mathcal{R}}$ as well, thus there exists $\mu_2 \in \mathcal{P}^{s_2,f}$ such that $[\mu_2]_{\mathcal{R}} = [\mu_1]_{\mathcal{R}}$. By definition of $[\cdot]_{\mathcal{R}}$, we have that for each $C \in S/\mathcal{R}$, $\sum_{s \in C} \mu_i(s) = [\mu_i]_{\mathcal{R}}$ for $i \in \{1,2\}$, thus $[\mu_2]_{\mathcal{R}} = [\mu_1]_{\mathcal{R}}$ implies that for each $C \in S/\mathcal{R}$, $\sum_{s \in C} \mu_1(s) = \sum_{s \in C} \mu_2(s)$, i.e., $\mu_1 \; \mathcal{L}(\mathcal{R}) \; \mu_2$. This means that we have found $s_2 \longrightarrow \mu_2$ with $\mu_1 \; \mathcal{L}(\mathcal{R}) \; \mu_2$, as required. □

By using Lemmas 2 and 5 and Proposition 1, we can now show that $\sim_c$ is preserved by the synchronous product operator introduced in Definition 10.

Theorem 2. *Given three IMDPs $\mathfrak{I}_1$, $\mathfrak{I}_2$, and $\mathfrak{I}_3$, if $\mathfrak{I}_1 \sim_c \mathfrak{I}_2$, then $\mathfrak{I}_1 \otimes \mathfrak{I}_3 \sim_c \mathfrak{I}_2 \otimes \mathfrak{I}_3$.*

Proof. Assume that $\mathfrak{I}_1 \sim_c \mathfrak{I}_2$. By Lemma 2, it follows that $\mathtt{UF}(\mathfrak{I}_1) \sim^p_{aa} \mathtt{UF}(\mathfrak{I}_2)$, thus, by Proposition 1, we have that $\mathtt{UF}(\mathfrak{I}_1) \otimes \mathtt{UF}(\mathfrak{I}_3) \sim^p_{aa} \mathtt{UF}(\mathfrak{I}_2) \otimes \mathtt{UF}(\mathfrak{I}_3)$. Lemma 5 now implies that $\mathtt{F}(\mathtt{UF}(\mathfrak{I}_1) \otimes \mathtt{UF}(\mathfrak{I}_3)) \sim_c \mathtt{F}(\mathtt{UF}(\mathfrak{I}_2) \otimes \mathtt{UF}(\mathfrak{I}_3))$, that is, $\mathfrak{I}_1 \otimes \mathfrak{I}_3 \sim_c \mathfrak{I}_2 \otimes \mathfrak{I}_3$, as required. □

5. Interleaved approach

In the previous sections, we have considered the parallel composition via synchronous production, which is working by the definition of folding collapsing all labels to a single transition. Here we consider the other extreme of the parallel composition: interleaving only.

Definition 11. Given two *IMDPs* $\mathfrak{I}_l$ and $\mathfrak{I}_r$, we define the *interleaved composition* of $\mathfrak{I}_l$ and $\mathfrak{I}_r$, denoted by $\mathfrak{I}_l \between \mathfrak{I}_r$, as the *IMDP* $\mathfrak{I} = (S, \bar{s}, \mathcal{A}, \mathtt{AP}, L, I)$ where $S = S_l \times S_r$; $\bar{s} = (\bar{s}_l, \bar{s}_r)$; $\mathcal{A} = (\mathcal{A}_l \times \{l\}) \cup (\mathcal{A}_r \times \{r\})$; $\mathtt{AP} = \mathtt{AP}_l \cup \mathtt{AP}_r$; for each $(s_l, s_r) \in S$, $L(s_l, s_r) = L_l(s_l) \cup L_r(s_r)$; and

$$I((s_l, s_r), (a,i), (t_l, t_r)) = \begin{cases} I_l(s_l, a, t_l) & \text{if } i = l \text{ and } t_r = s_r, \\ I_r(s_r, a, t_r) & \text{if } i = r \text{ and } t_l = s_l, \\ [0,0] & \text{otherwise.} \end{cases}$$

Theorem 3. *Given three IMDPs $\mathfrak{I}_1$, $\mathfrak{I}_2$, and $\mathfrak{I}_3$, if $\mathfrak{I}_1 \sim_c \mathfrak{I}_2$, then $\mathfrak{I}_1 \between \mathfrak{I}_3 \sim_c \mathfrak{I}_2 \between \mathfrak{I}_3$.*

Proof. Let $\mathcal{R}$ be the probabilistic bisimulation justifying $\mathfrak{I}_1 \sim_c \mathfrak{I}_2$ and define $\mathcal{R}' = \mathcal{R} \times \mathcal{I}_{S_3}$; we claim that $\mathcal{R}'$ is a probabilistic bisimulation between $\mathfrak{I}_1 \between \mathfrak{I}_3$ and $\mathfrak{I}_2 \between \mathfrak{I}_3$. The fact that $\mathcal{R}'$ is an equivalence relation follows trivially by its definition and the fact that $\mathcal{R}$ is an equivalence relation. The fact that $((\bar{s}_1, \bar{s}_3), (\bar{s}_2, \bar{s}_3))$ follows immediately by the hypothesis that $(\bar{s}_1, \bar{s}_2) \in \mathcal{R}$ and $(\bar{s}_3, \bar{s}_3) \in \mathcal{I}_{S_3}$.

Let $((s_1, s_3), (s_2, s_3)) \in \mathcal{R}'$. Assume, without loss of generality, that $s_1 \in S_1$ and $s_2 \in S_2$; the other cases are similar. The fact that $L_{1,3}(s_1, s_3) = L_{2,3}(s_2, s_3)$ is straightforward, since by definition of interleaved composition and the hypothesis that $s_1 \; \mathcal{R} \; s_2$, we have that $L_{1,3}(s_1, s_3) = L_1(s_1) \cup L_3(s_3) = L_2(s_2) \cup L_3(s_3) = L_{2,3}(s_2, s_3)$, as required.

Consider now a transition $(s_1, s_3) \longrightarrow \mu_{1,3}$. By definition, we have that $\mu_{1,3} \in \mathrm{CH}(\bigcup_{(a,i) \in \mathcal{A}(s_1,s_3)} \mathcal{P}^{(s_1,s_3),(a,i)})$. This implies that there exist a multiset of distributions $\{ \mu_{a,i} \in \mathcal{P}^{(s_1,s_3),(a,i)} \mid (a,i) \in$

$\mathcal{A}(s_1, s_3)\}$ and a multiset of real values $\{\, p_{a,i} \in \mathbb{R}_{\geq 0} \mid (a,i) \in \mathcal{A}(s_1, s_3)\,\}$ such that $\sum_{(a,i)\in\mathcal{A}(s_1,s_3)} p_{a,i} = 1$ and $\sum_{(a,i)\in\mathcal{A}(s_1,s_3)} p_{a,i} \cdot \mu_{a,i} = \mu_{1,3}$. Consider an action $(a,i) \in \mathcal{A}(s_1, s_3)$: by definition of interleaved composition, it is either of the form $(a,l) \in \mathcal{A}_l \times \{l\}$, or of the form $(a,r) \in \mathcal{A}_r \times \{r\}$. Consider the two cases separately:

Case $(a,i) \in \mathcal{A}_l \times \{l\}$**:** this means that $\mu_{a,i}$ is actually the distribution $\mu_{a,i} = \mu_a \times \delta_{s_3}$ where $\mu_a \in \mathcal{P}^{s_1,a}$ is such that for each $s_1' \in S_1$, $\mu_a(s_1') = \mu_{a,i}(s_1', s_3)$, thus $s_1 \longrightarrow \mu_a$. Since by hypothesis $(s_1, s_2) \in \mathcal{R}$ and $\mathcal{R}$ is a probabilistic bisimulation, there exists $\mu_{a,2} \in \mathrm{CH}(\bigcup_{b\in\mathcal{A}(s_2)} \mathcal{P}^{s_2,b})$ such that $\mu_a \; \mathcal{L}(\mathcal{R}) \; \mu_{a,2}$. This implies that there exist a multiset of distributions $\{\mu_{a,2,b} \in \mathcal{P}^{s_2,b} \mid b \in \mathcal{A}(s_2)\}$ and a multiset of real values $\{\, p_{a,2,b} \mid b \in \mathcal{A}(s_2)\,\}$ such that $\sum_{b\in\mathcal{A}(s_2)} p_{a,2,b} = 1$, $\sum_{b\in\mathcal{A}(s_2)} p_{a,2,b} \cdot \mu_{a,2,b} = \mu_{a,2}$, and $\mu_a \times \delta_{s_3} \; \mathcal{L}(\mathcal{R}') \; \mu_{a,2} \times \delta_{s_3}$. This means that for each $b \in \mathcal{A}(s_2)$, we have that $\mu_{a,2,b} \times \delta_{s_3} \in \mathcal{P}^{(s_2,s_3),(b,l)}$, thus by taking $\mu_{a,2,3} = \sum_{b\in\mathcal{A}(s_2)} p_{a,2,b} \cdot (\mu_{a,2,b} \times \delta_{s_3})$ we have that $(s_2, s_3) \longrightarrow \mu_{a,2,3}$ and $\mu_{a,i} \; \mathcal{L}(\mathcal{R}') \; \mu_{a,2,3}$.

Case $(a,i) \in \mathcal{A}_r \times \{r\}$**:** this means that $\mu_{a,i}$ is actually the distribution $\mu_{a,i} = \delta_{s_1} \times \mu_a$ where $\mu_a \in \mathcal{P}^{s_3,a}$ is such that for each $s_3' \in S_3$, $\mu_a(s_3') = \mu_{a,i}(s_1, s_3')$, thus $s_3 \longrightarrow \mu_a$. This implies trivially that $(s_2, s_3) \longrightarrow \mu_{a,2,3}$ where $\mu_{a,2,3} = \delta_{s_2} \times \mu_a$ and $\mu_{a,i} \; \mathcal{L}(\mathcal{R}') \; \mu_{a,2,3}$.

From the analysis of the two cases, we have that for each $(a,i) \in \mathcal{A}(s_1, s_3)$, there exists a transition $(s_2, s_3) \longrightarrow \mu_{a,2,3}$ such that $\mu_{a,i} \; \mathcal{L}(\mathcal{R}') \; \mu_{a,2,3}$. This implies that $\mu_{2,3} = \sum_{(a,i)\in\mathcal{A}(s_1,s_3)} p_{a,i} \cdot \mu_{a,2,3} \in \mathrm{CH}(\bigcup_{(b,j)\in\mathcal{A}(s_2,s_3)} \mathcal{P}^{(s_2,s_3),(b,j)})$ and $\mu_{1,3} \; \mathcal{L}(\mathcal{R}') \; \mu_{2,3}$, as required. □

6. Concluding Remarks

In this paper, we have studied the probabilistic bisimulation problem for interval *MDPs* in order to speed up the run time of model checking algorithms that often suffer from the state space explosion. Interval MDPs include two sources of nondeterminism for which we have considered the cooperative resolution in a dynamic setting. We have revised and extended the compositionality reasoning in [9] by further exploration on the possibility of defining the parallel operator for *IMDP* models which preserve our notion of probabilistic bisimulation.

Acknowledgments. This work is supported by the EU 7th Framework Programme under grant agreements 295261 (MEALS) and 318490 (SENSATION), by the DFG as part of SFB/TR 14 AVACS, by the ERC Advanced Investigators Grant 695614 (POWVER), by the CAS/SAFEA International Partnership Program for Creative Research Teams, by the National Natural Science Foundation of China (Grants No. 61472473, 61532019, 61550110249, 61550110506), by the Chinese Academy of Sciences Fellowship for International Young Scientists, and by the CDZ project CAP (GZ 1023).

References

[1] M. Benedikt, R. Lenhardt, and J. Worrell. LTL model checking of interval Markov chains. In *TACAS*, volume 7795 of *LNCS*, pages 32–46, 2013.

[2] S. Cattani and R. Segala. Decision algorithms for probabilistic bisimulation. In *CONCUR*, volume 2421 of *LNCS*, pages 371–385, 2002.

[3] K. Chatterjee, K. Sen, and T. A. Henzinger. Model-checking omega-regular properties of interval Markov chains. In *FoSSaCS*, volume 4962 of *LNCS*, pages 302–317, 2008.

[4] G. Chehaibar, H. Garavel, L. Mounier, N. Tawbi, and F. Zulian. Specification and verification of the PowerScale® bus arbitration protocol: An industrial experiment with LOTOS. In *FORTE*, pages 435–450, 1996.

[5] N. Coste, H. Hermanns, E. Lantreibecq, and W. Serwe. Towards performance prediction of compositional models in industrial GALS designs. In *CAV*, volume 5643 of *LNCS*, pages 204–218, 2009.

[6] B. Delahaye, J.-P. Katoen, K. G. Larsen, A. Legay, M. L. Pedersen, F. Sher, and A. Wasowski. Abstract probabilistic automata. In *VMCAI*, volume 6538 of *LNCS*, pages 324–339, 2011.

[7] B. Delahaye, K. G. Larsen, A. Legay, M. L. Pedersen, and A. Wasowski. Decision problems for interval Markov chains. In *LATA*, volume 6638 of *LNCS*, pages 274–285, 2011.

[8] V. Hashemi, H. Hatefi, and J. Krčál. Probabilistic bisimulations for PCTL model checking of interval MDPs. In *SynCoP*, volume 145 of *EPTCS*, pages 19–33, 2014.

[9] V. Hashemi, H. Hermanns, L. Song, K. Subramani, A. Turrini, and P. Wojciechowski. Compositional bisimulation minimization for interval Markov decision processes. In *LATA*, volume 9618 of *LNCS*, pages 114–126, 2016.

[10] H. Hermanns. *Interactive Markov chains: and the quest for quantified quality*. Springer-Verlag, 2002.

[11] H. Hermanns and J.-P. Katoen. Automated compositional Markov chain generation for a plain-old telephone system. *Sci. Comp. Progr.*, 36(1):97–127, 2000.

[12] G. N. Iyengar. Robust dynamic programming. *Mathematics of Operations Research*, 30(2):257–280, 2005.

[13] B. Jonsson and K. G. Larsen. Specification and refinement of probabilistic processes. In *LICS*, pages 266–277, 1991.

[14] P. C. Kanellakis and S. A. Smolka. CCS expressions, finite state processes, and three problems of equivalence. *Information and Computation*, 86(1):43–68, 1990.

[15] J.-P. Katoen, T. Kemna, I. S. Zapreev, and D. N. Jansen. Bisimulation minimisation mostly speeds up probabilistic model checking. In *TACAS*, volume 4424 of *LNCS*, pages 76–92, 2007.

[16] J.-P. Katoen, D. Klink, and M. R. Neuhäußer. Compositional abstraction for stochastic systems. In *FORMATS*, volume 5813 of *LNCS*, pages 195–211, 2009.

[17] I. Kozine and L. V. Utkin. Interval-valued finite Markov chains. *Reliable Computing*, 8(2):97–113, 2002.

[18] M. Kwiatkowska, G. J. Norman, and D. Parker. PRISM 4.0: Verification of probabilistic real-time systems. In *CAV*, volume 6806 of *LNCS*, pages 585–591, 2011.

[19] R. Paige and R. E. Tarjan. Three partition refinement algorithms. *SIAM Journal on Computing*, 16(6):973–989, 1987.

[20] R. Segala. *Modeling and Verification of Randomized Distributed Real-Time Systems*. PhD thesis, MIT, 1995.

[21] R. Segala. Probability and nondeterminism in operational models of concurrency. In *CONCUR*, volume 4137 of *LNCS*, pages 64–78, 2006.

[22] K. Sen, M. Viswanathan, and G. Agha. Model-checking Markov chains in the presence of uncertainties. In *TACAS*, volume 3920 of *LNCS*, pages 394–410, 2006.

[23] A. Turrini and H. Hermanns. Cost preserving bisimulations for probabilistic automata. *Logical Methods in Computer Science*, 4(11):1–58, 2014.

[24] W. Yi. Algebraic reasoning for real-time probabilistic processes with uncertain information. In *FTRTFT*, volume 863 of *LNCS*, pages 680–693, 1994.